

CLAIM AMENDMENTS

Please amend the claims as indicated:

1 1. (currently amended) ~~A computer system comprising~~ In a computer system
2 that includes:
3 at least one hardware processor; and
4 a first operating system (COS) initially installed to run on the hardware processor
5 at a most-privileged, system level, the system level being defined as an operational
6 state with permission to directly access predetermined physical resources of the
7 computer,
8 a method comprising:
9 initializing the computer with the COS;
10 ~~the COS forming means for initializing the computer;~~
11 ~~a kernel that forms a second operating system;~~
12 ~~loading means for loading the kernel via the COS, loading and for starting~~
13 execution of the a kernel, the kernel thereupon substantially displacing the COS from
14 the system level and itself running at the system level; and
15 ~~the kernel forming means for handling requests for system resources via the~~
16 kernel, including scheduling execution of the COS on the hardware processor(s).

1 2. (original) A method as in claim 1, in which the step of loading the kernel
2 comprises:
3 loading a load call module within the COS;
4 upon initialization of the computer, calling a loading module from the load call
5 module, whereupon the loading module loads the kernel.

1 3. (original) A method as defined in claim 2, in which the step of loading the
2 load call module within the COS comprises installing the load call module as a driver
3 within the COS.

1 4. (currently amended) In a computer system that includes:
2 at least one hardware processor that has a hardware instruction pointer; and
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer;

7 a method comprising:
8 initializing the computer with the COS;
9 loading a load call module within the COS;
10 upon initialization of the computer, calling a loading module from the load call
11 module, whereupon the loading module loads a kernel, which forms a second operating
12 system;

13 via the COS starting execution of the kernel, the kernel thereupon substantially
14 displacing the COS from the system level and itself running at the system level;

15 ~~A method as in claim 2, in which the computer includes at least one processor,~~
16 ~~which has a hardware instruction pointer, and the step of loading the kernel including~~
17 ~~the following sub-steps:~~

18 via the loading module, setting the hardware instruction pointer and forwarding of
19 interrupts and faults generated by the processor and by predetermined ones of the
20 system resources to point into a memory address space allocated to and controlled by
21 the kernel, the kernel thereby handling requests for system resources.

1 5. (original) A method as in claim 4, further including the following steps:

2 after initialization of the computer, transferring from the COS to the kernel a list
3 of devices initially controlled by the COS, the devices being included among the system
4 resources; and

5 classifying the devices and control of the devices into the following groups (which
6 may be empty):

7 host-managed devices, which are controlled by the COS;

8 reserved devices, which are controlled by the kernel;

9 and shared devices, which may be controlled by either the COS or the

10 kernel.

1 6. (original) A method as defined in claim 5, further comprising including a
2 mass storage controller as one of the shared devices.

1 7. (original) A method as defined in claim 6, in which the mass storage
2 controller is a SCSI device.

1 8. (original) A method as defined in claim 5, further comprising including a
2 network adapter as one of the shared devices.

1 9. (original) A method as defined in claim 5, further comprising the steps of
2 forwarding interrupts generated by host-managed devices to the COS via the kernel,
3 and handling such interrupts within the COS.

1 10. (original) A method as defined in claim 9, further including the step of
2 delaying handling of interrupts that are forwarded to the COS and that are generated by
3 host-managed devices until a subsequent instance of running of the COS.

1 11. (original) A method as defined in claim 10, further including the step, upon
2 sensing, in the kernel, an interrupt raised by any host-managed device, of masking the
3 interrupt until the subsequent instance of running of the COS, thereby avoiding multiple
4 recurrences of the interrupt.

1 12. (original) A method as defined in claim 1, further including the step of
2 installing at least one virtual machine (VM) to run on the kernel via a virtual machine
3 monitor (VMM).

1 13. (currently amended) In a computer system that includes:
2 at least one hardware processor;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer;
7 a method comprising:
8 initializing the computer with the COS;
9 via the COS, loading and starting execution of a kernel, which forms a second
10 operating system, the kernel thereupon substantially displacing the COS from the
11 system level and itself running at the system level, the kernel thereby handling requests
12 for system resources;
13 ~~A method as defined in claim 12, further including the following steps:~~
14 installing at least one virtual machine (VM) to run on the kernel via a virtual
15 machine monitor (VMM);
16 in the kernel, separately scheduling the execution of the COS and of each VM,
17 the COS and the ~~VM's~~ VMs thereby forming separately schedulable and separately
18 executing entities; and
19 within the kernel, representing each schedulable entity as a corresponding world,
20 each world comprising a world memory region with a respective world address space
21 and storing a respective world control thread.

1 14. (original) A method as defined in claim 13, further including the step of
2 switching worlds, which step comprises:
3 under control of the kernel, storing current state data for a currently executing
4 schedulable entity in a kernel-controlled memory region;
5 disabling exceptions;
6 loading state data for a subsequently executing schedulable entity;
7 starting execution of the subsequently executing schedulable entity; and
8 enabling exceptions.

1 15. (original) A method as defined in claim 14, in which the state data for each
2 schedulable entity includes exception flags, memory segments, an instruction pointer,
3 and descriptor tables, which are loaded into an exception flags register, memory
4 segment registers, an instruction pointer register, and descriptor tables, respectively.

1 16. (original) A method as defined in claim 14, in which the computer includes
2 a plurality of hardware processors, further including the following steps:
3 in the kernel, separately scheduling the execution of each processor, the
4 processors thereby also forming separately schedulable entities;
5 within the kernel, representing each processor as a corresponding system world,
6 each system having a respective system world address space and a respective system
7 world control thread.

1 17. (original) A method as defined in claim 16, further including the step of
2 allocating, for each processor, a separate memory mapping cache.

1 18. (original) A method as defined in claim 13, in which each VM includes a
2 virtual processor, a virtual operating system (VOS), and an I/O driver, loaded within the
3 VOS, for an I/O device, the method further comprising the following steps:
4 allocating a shared memory space that is addressable by both the kernel and the
5 VM's I/O driver,
6 transferring an output set of data from the VM to the I/O device according to the
7 following sub-steps:
8 via the VM's I/O driver, setting a pointer to the output set of data in the
9 shared memory region and generating a request for transmission;
10 in the kernel, upon sensing the request for transmission:
11 retrieving the output set of data from a position in the shared memory region
12 indicated by the pointer and transferring the retrieved output set of data to a physical
13 transmit buffer portion of the shared memory region;
14 transferring the output data set from the physical transmit buffer portion to the
15 I/O device;

16 transferring an input set of data from the I/O device to the VM according to the
17 following sub-steps:
18 in the kernel,
19 copying the input set of data into a physical receive buffer portion
20 of the shared memory region;
21 setting the pointer to the physical receive buffer portion;
22 issuing to the VMM an instruction to raise an interrupt;
23 in the VM, upon sensing the interrupt raised by the VMM, retrieving the
24 input set of data from the physical receive buffer portion of the shared memory region
25 indicated by the pointer;
26 whereby the input and output data sets may be transferred between the VM and
27 the I/O device via only the kernel.

1 19. (original) A method as defined in claim 18, further comprising completing
2 the sub-steps for transferring the output set of data upon sensing only a single request
3 for transmission.

1 20. (original) A method as defined in claim 18, in which:
2 the I/O device is a network connection device for data transfer to and from a
3 network; and
4 the input and output data sets are network packets.

1 21. (original) A method as defined in claim 12, further including the following
2 steps:
3 mapping a kernel address space, within which the kernel is stored and is
4 addressable by the kernel, into a VMM address space, within which the VMM is stored
5 and which is addressable by the VMM.

1 22. (original) A method as defined in claim 21, in which the computer has a
2 segmented memory architecture, the memory being addressable via segment registers,
3 further including the step of setting a segment length for the VMM larger than a
4 minimum length necessary to fully contain both the VMM and the kernel, whereby the
5 step of mapping the kernel address space within the VMM address space may be
6 performed free of any need to change a corresponding segment register.

1 23. (original) A method as defined in claim 12, in which each VM includes a
2 virtual processor, a virtual operating system (VOS), and a virtual disk (VDISK), the
3 method further including carrying out the following steps within the kernel:
4 partitioning the VDISK into VDISK blocks;
5 maintaining an array of VDISK block pointers, the array comprising a plurality of
6 sets of VDISK block pointers;
7 maintaining a file descriptor table in which is stored file descriptors, each file
8 descriptor storing block identification and allocation information, and at least one pointer
9 block pointer;
10 each pointer block pointer pointing to one of the sets of VDISK block pointers;
11 and
12 each VDISK block pointer identifying the location of a respective one of the
13 VDISK blocks.

1 24. (original) A method as defined in claim 1, further including the following
2 steps:
3 halting execution of the kernel;
4 reinstating a state of the first operating system that existed before the loading of
5 the kernel; and
6 resuming execution of the first operating system at the most-privileged system
7 level;
8 the kernel thereby being functionally removed from the computer.

1 25. (currently amended) In a computer system that includes:
2 at least one hardware processor, which has a hardware instruction pointer;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer;
7 a method comprising:
8 initializing the computer with the COS;
9 via the COS, loading and starting execution of a kernel, which forms a second
10 operating system, the kernel thereupon substantially displacing the COS from the
11 system level and itself running at the system level, the kernel thereby handling requests
12 for system resources;
13 halting execution of the kernel;
14 reinstating a state of the first operating system that existed before the loading of
15 the kernel; and
16 resuming execution of the first operating system at the most-privileged system
17 level;
18 the kernel thereby being functionally removed from the computer;
19 ~~A method as defined in claim 24, in which:~~
20 ~~A) the computer includes at least one processor, which has a hardware~~
21 ~~instruction pointer;~~
22 B A) the step of loading the kernel includes the following sub-steps:
23 i) loading a load call module within the COS;
24 ii) upon initialization of the computer, calling a loading module from the
25 load call module, whereupon the loading module loads the kernel;
26 iii) after initialization of the computer, transferring from the COS to the
27 kernel a list of devices initially controlled by the COS; and
28 iv) classifying the devices and control of the devices into the following
29 groups (which may be empty):
30 a) host-managed devices, which are controlled by the COS;
31 b) reserved devices, which are controlled by the kernel; and

32 c) shared devices, which may be controlled by either the COS or
33 the kernel;

34 v) via the loading module, setting the hardware instruction pointer and
35 forwarding of interrupts and faults generated by the processor and by predetermined
36 ones of the physical resources to point into a memory address space allocated to and
37 controlled by the kernel;

38 ¶ B) the step of reinstating the state of the first operating system includes the
39 following steps:

40 i) restoring interrupt and fault handling from the kernel to the first
41 operating system;

42 ii) transferring control of host-managed and shared devices from the
43 kernel to the first operating system; and

44 iii) removing the kernel from an address space of the first operating
45 system.

1 26. (Original) A method for managing resources in a computer, which includes
2 at least one processor that has a hardware instruction pointer, the method comprising
3 the following steps:

4 A) initializing the computer using a first operating system (COS), the COS itself
5 running at a most-privileged, system level, the system level being defined as an
6 operational state with permission to directly access predetermined physical resources
7 of the computer, the physical resources including physical devices;

8 B) loading a kernel via the COS, the kernel forming a second operating system,
9 this step of loading the kernel comprising:

10 i) loading a load call module within the COS;

11 ii) upon initialization of the computer, calling a loading module from the
12 load call module, whereupon the loading module loads the kernel;

13 iii) via the loading module, setting the hardware instruction pointer and
14 forwarding interrupts and faults generated by the processor and by predetermined ones
15 of the physical resources to point into a memory address space allocated to and
16 controlled by the kernel;

17 C) starting execution of the kernel, the kernel thereupon substantially displacing
18 the COS from the system level and itself running at the system level; and
19 D) submitting requests for system resources via the kernel;
20 E) after initialization of the computer, transferring from the COS to the kernel a
21 list of the devices initially controlled by the COS; and
22 F) classifying the devices and control of the devices into the following groups
23 (which may be empty):
24 i) host-managed devices, which are controlled by the COS;
25 ii) reserved devices, which are controlled by the kernel; and
26 iii) shared devices, which may be controlled by either the COS or the
27 kernel; and
28 G) forwarding interrupts generated by host-managed devices to the COS via the
29 kernel, and handling such interrupts within the COS.

1 27. (currently amended) A method for managing resources in a computer,
2 which includes at least one processor that has a hardware instruction pointer, the
3 method comprising the following steps:

4 A) initializing the computer using a first operating system (COS), the COS itself
5 running at a most-privileged, system level, the system level being defined as an
6 operational state with permission to directly access predetermined physical resources
7 of the computer, the physical resources including physical devices;

8 B) loading a kernel via the COS, the kernel forming a second operating system,
9 this step of loading the kernel comprising:

10 i) loading a load call module within the COS;
11 ii) upon initialization of the computer, calling a loading module from the
12 load call module, whereupon the loading module loads the kernel;
13 iii) via the loading module, setting the hardware instruction pointer and
14 forwarding interrupts and faults generated by the processor and by predetermined ones
15 of the physical resources to point into a memory address space allocated to and
16 controlled by the kernel;

17 C) starting execution of the kernel, the kernel thereupon substantially displacing
18 the COS from the system level and itself running at the system level; and

19 D) submitting requests for system resources via the kernel;
20 E) after initialization of the computer, transferring from the COS to the kernel a
21 list of the devices initially controlled by the COS; and
22 F) classifying the devices and control of the devices into the following groups
23 (which may be empty):
24 i) host-managed devices, which are controlled by the COS;
25 ii) reserved devices, which are controlled by the kernel; and
26 iii) shared devices, which may be controlled by either the COS or the
27 kernel; and
28 G) forwarding interrupts generated by host-managed devices to the COS via the
29 kernel, and handling such interrupts within the COS;
30 H) installing at least one virtual machine (VM) to run on the kernel via a virtual
31 machine monitor (VMM); and
32 I) in the kernel, separately scheduling the execution of the COS and of each
33 VM, the COS and the ~~VM's~~ VMs thereby forming separately schedulable and separately
34 executing entities.

1 28. (currently amended) A computer system comprising:
2 at least one hardware processor;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer, the COS forming means for initializing the computer;
7 a kernel ~~means~~ that forms a second operating system;
8 ~~a loading means~~ loader comprising computer-executable code for loading the
9 kernel means via the COS and for starting execution of the kernel means, the kernel
10 ~~means~~ thereupon substantially displacing the COS from the system level and itself
11 running at the system level; and
12 the kernel ~~means is provided~~ including a software module for handling requests
13 for system resources, including scheduling execution of the COS on the hardware
14 processor(s).

1 29. (currently amended) A system as in claim 28, in which the ~~loading means~~
2 loader includes a loading driver installed within the COS.

1 30. (currently amended) A computer system comprising:
2 at least one hardware processor that has a hardware instruction pointer;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer, the COS initializing the computer;
7 a kernel that forms a second operating system;
8 a loader comprising computer-executable code for loading the kernel via the
9 COS and for starting execution of the kernel, the kernel thereupon substantially
10 displacing the COS from the system level and itself running at the system level;
11 the kernel including a software module for handling requests for system
12 resources;

13 ~~A system as in claim 28, in which:~~
14 ~~the processor has a hardware instruction pointer;~~
15 ~~the loading means is loader being~~ further provided for setting the hardware
16 instruction pointer and forwarding interrupts and faults generated by the processor and
17 by predetermined ones of the system resources to point into a memory address space
18 allocated to and controlled by the kernel ~~means~~.

1 31. (currently amended) A system as in claim 30, in which:
2 the system resources include devices initially controlled by the COS;
3 ~~the loading means loader~~ is further provided for transferring, after initialization of
4 the computer, from the COS to the kernel ~~means~~ a list of the devices initially controlled
5 by the COS and for classifying the devices and control of the devices into the following
6 groups (which may be empty):
7 host-managed devices, which are controlled by the COS;
8 reserved devices, which are controlled by the kernel ~~means~~; and
9 shared devices, which may be controlled by either the COS or the kernel
10 ~~means~~.

1 32. (original) A system as in claim 31, in which at least one of the shared
2 devices is a mass storage controller.

1 33. (currently amended) A system as defined in claim 28, further comprising:
2 at least one virtual machine (VM); and
3 a virtual machine monitor (VMM);
4 in which the VM is installed to run on the kernel means via the VMM.

1 34. (currently amended) ~~A system as defined in claim 33,~~ A computer system
2 comprising:
3 at least one hardware processor;
4 a first operating system (COS) initially installed to run on the hardware processor
5 at a most-privileged, system level, the system level being defined as an operational
6 state with permission to directly access predetermined physical resources of the
7 computer, the COS forming means for initializing the computer;
8 a kernel that forms a second operating system;
9 a loader comprising computer-executable code for loading the kernel means via
10 the COS and for starting execution of the kernel means, the kernel thereupon
11 substantially displacing the COS from the system level and itself running at the system
12 level;
13 a virtual machine monitor (VMM);
14 at least one virtual machine (VM) installed to run on the kernel via the VMM;
15 the kernel including software modules
16 for handling requests for system resources;
17 ~~in which the kernel means is further provided:~~
18 for separately scheduling the execution of the COS and of each VM, the
19 COS and the ~~VM's~~ VMs thereby forming separately schedulable and separately
20 executing entities; and
21 for representing each schedulable entity as a corresponding world, each
22 world comprising a world memory region with a respective world address space and
23 storing a respective world control thread.

1 35. (currently amended) A system as defined in claim 34, in which the kernel
2 means is further provided:

3 for storing current state data for a currently executing schedulable entity in a
4 kernel-controlled memory region;

5 for disabling exceptions;

6 for loading state data for a subsequently executing schedulable entity;

7 for starting execution of the subsequently executing schedulable entity; and

8 for enabling exceptions;

9 the kernel means thereby being provided for switching worlds.

1 36. (original) A system as defined in claim 35, in which the state data for each
2 schedulable entity includes exception flags, memory segments, an instruction pointer,
3 and descriptor tables, which are loaded into an exception flags register, memory
4 segment registers, an instruction pointer register, and descriptor tables, respectively.

1 37. (currently amended) A system as defined in claim 34, in which:
2 the computer includes a plurality of hardware processors;
3 the kernel means is further provided:
4 for separately scheduling the execution of each processor, the processors
5 thereby also forming separately schedulable entities;
6 for representing each processor as a corresponding system world, each
7 system having a respective system world address space and a respective system world
8 control thread.

1 38. (original) A system as defined in claim 37, further comprising a separate
2 memory mapping cache for each processor.

1 39. (currently amended) A computer system as defined in claim 33, further
2 comprising:
3 at least one hardware processor;
4 a first operating system (COS) initially installed to run on the hardware processor
5 at a most-privileged, system level, the system level being defined as an operational
6 state with permission to directly access predetermined physical resources of the
7 computer, the COS forming means for initializing the computer;
8 a kernel that forms a second operating system;
9 a loader comprising computer-executable code for loading the kernel means via
10 the COS and for starting execution of the kernel means, the kernel thereupon
11 substantially displacing the COS from the system level and itself running at the system
12 level;
13 the kernel handling requests for system resources;
14 a virtual machine monitor (VMM);
15 at least one virtual machine (VM) installed to run on the kernel via the VMM;
16 within each VM, a virtual processor, a virtual operating system (VOS), and an I/O
17 driver for an I/O device loaded within the VOS;
18 a shared memory space that is addressable by both the kernel means and the
19 VM's I/O driver, the shared memory space storing input data and output data for
20 transfer between the VM and the I/O device;
21 in which:
22 the VM's I/O driver ~~forms means~~ comprises computer-executable code for
23 setting a pointer to output data in the shared memory region and generating a request
24 for transmission;
25 the kernel ~~means~~ is further provided, upon sensing the request for transmission:
26 for retrieving the output data from a position in the shared memory region
27 indicated by the pointer and transferring the retrieved output data to a physical transmit
28 buffer portion of the shared memory region;
29 for outputting the output data from the physical transmit buffer portion to
30 the I/O device;
31 for receiving the input data from the I/O device;

32 for copying the input data into a physical receive buffer portion of the
33 shared memory region;
34 for setting the pointer to the physical receive buffer portion;
35 for issuing to the VMM an instruction to raise an interrupt;
36 the VM's I/O driver ~~forming means~~ is further provided, upon sensing the interrupt
37 raised by the VMM, for retrieving the input data from the physical receive buffer portion
38 of the shared memory region indicated by the pointer;
39 whereby the input and output data may be transferred between the VM and the
40 I/O device via only the kernel ~~means~~.

1 40. (original) A system as defined in claim 39, in which:
2 the I/O device is a network connection device for data transfer to and from a
3 network; and
4 the input and output data are network packets.

1 41. (currently amended) A system as defined in claim 33, further comprising:
2 a kernel address memory portion, within which the kernel ~~means~~ is stored and is
3 addressable by the kernel ~~means~~;
4 ~~means~~ a mapping module comprising computer-executable code for mapping
5 the kernel address memory portion into a VMM address space, within which the VMM is
6 stored and which is addressable by the VMM.

1 42. (currently amended) A system as defined in claim 41, in which:
2 the computer has a segmented memory architecture;
3 segment registers via which the memory is addressed; and
4 a segment length for the VMM that is larger than a minimum length necessary to
5 fully contain both the VMM and the kernel ~~means~~.

1 43. (currently amended) A system as defined in claim 33, in which:
2 each VM includes a virtual processor, a virtual operating system (VOS), and a
3 virtual disk (VDISK);
4 the VDISK is partitioning into VDISK blocks;
5 the kernel means is further provided:
6 for maintaining an array of VDISK block pointers, the array comprising a
7 plurality of sets of VDISK block pointers;
8 for maintaining a file descriptor table in which is stored file descriptors,
9 each file descriptor storing block identification and allocation information, and at least
10 one pointer block pointer;
11 each pointer block pointer pointing to one of the sets of VDISK block pointers;
12 and
13 each VDISK block pointer identifying the location of a respective one of the
14 VDISK block.

1 44. (original) A computer system comprising:
2 at least one hardware processor;
3 a first operating system (COS) initially installed to run on the hardware processor
4 at a most-privileged, system level, the system level being defined as an operational
5 state with permission to directly access predetermined physical resources of the
6 computer, the COS forming means for initializing the computer;
7 a kernel means that forms a second operating system;
8 loading means for loading the kernel via the COS and for starting execution of
9 the kernel, the kernel thereupon substantially displacing the COS from the system level
10 and itself running at the system level;
11 the kernel means is provided for handling requests for system resources;
12 in which:
13 the processor has a hardware instruction pointer;
14 the loading means is further provided for setting the hardware instruction pointer
15 and forwarding interrupts and faults generated by the processor and by predetermined
16 ones of the system resources to point into a memory address space allocated to and
17 controlled by the kernel means;

the system resources include devices initially controlled by the COS;
the loading means is further provided for transferring, after initialization of the computer, from the COS to the kernel means a list of the devices initially controlled by the COS and for classifying the devices and control of the devices into the following groups (which may be empty):
host-managed devices, which are controlled by the COS;
reserved devices, which are controlled by the kernel means; and
shared devices, which may be controlled by either the COS or the kernel means.

45. (currently amended) A computer system comprising:
at least one hardware processor;
a first operating system (COS) initially installed to run on the hardware processor at a most-privileged, system level, the system level being defined as an operational state with permission to directly access predetermined system resources of the computer, the COS forming means for initializing the computer;
a kernel means that forms a second operating system;
loading means for loading the kernel means via the COS and for starting execution of the kernel means, the kernel means thereupon substantially displacing the COS from the system level and itself running at the system level;
at least one virtual machine (VM); and
a virtual machine monitor (VMM);
in which:
the VM is installed to run on the kernel means via the VMM;
the kernel means is provided for handling requests for system resources;
the processor has a hardware instruction pointer;
the loading means is further provided for setting the hardware instruction pointer and forwarding interrupts and faults generated by the processor and by predetermined ones of the system resources to point into a memory address space allocated to and controlled by the kernel means;
the system resources include devices initially controlled by the COS;

the loading means is further provided for transferring, after initialization of the computer, from the COS to the kernel means a list of the devices initially controlled by the COS and for classifying the devices and control of the devices into the following groups (which may be empty):

- host-managed devices, which are controlled by the COS;
- reserved devices, which are controlled by the kernel means; and
- shared devices, which may be controlled by either the COS or the kernel means;

for separately scheduling the execution of the COS and of each VM, the COS and the ~~VM's~~ VMs thereby forming separately schedulable and separately executing entities; and

the kernel means is further provided:

- for representing each schedulable entity as a corresponding world, each world comprising a world memory region with a respective world address space and storing a respective world control thread;

- for storing current state data for a currently executing schedulable entity in a kernel means-controlled memory region;

- for disabling exceptions;

- for loading state data for a subsequently executing schedulable entity;

- for starting execution of the subsequently executing schedulable entity;

and

- for enabling exceptions;

the kernel means thereby being provided for switching worlds.